3.5 GHz Federal Incumbent Protection Algorithms

Michael R. Souryal and Thao T. Nguyen
Communications Technology Laboratory
National Institute of Standards and Technology
Gaithersburg, Maryland, U.S.
{souryal, thao.t.nguyen}@nist.gov

Nickolas J. LaSorte
Office of Spectrum Management
National Telecommunications and Information Administration
Washington, D.C.
nlasorte@ntia.doc.gov

***Abstract*—The current standardized algorithm for U.S. federal incumbent protection in the 3.5 GHz band is sub-optimum in that it moves more commercial transmissions out of the protected channel than necessary, in most cases. This paper proposes a more efficient algorithm that packs transmissions in the allowed interference budget jointly across incumbent receiver azimuths rather than independently, resulting in 18% to 25% fewer commercial transmissions affected.**

I. Introduction

Regulatory rules for the Citizens Broadband Radio Service (CBRS) in the U.S. [1] permit commercial broadband users to operate in the radio frequency (RF) spectrum from 3550 MHz to 3700 MHz (3.5 GHz band) provided they do not compromise the operations of federal incumbents in and adjacent to this band. Industry standards for CBRS systems specify the mechanism for protecting federal incumbents from harmful interference [2]. This mechanism requires that Spectrum Access Systems (SASs)—centralized frequency coordinators that authorize access to the band by CBRS devices (CBSDs)—coordinate their authorizations using a common, standardized algorithm. Using agreed upon RF propagation and aggregate interference models, the algorithm identifies which authorized transmissions can continue in a protected channel and which must be suspended or possibly relocated to another channel. The list of transmissions that must be suspended or relocated to protect a given channel occupied by the incumbent is referred to as the "move list" for that channel.

This paper is a study of alternative move-list algorithms for federal incumbent protection. We show that in most cases the current standardized algorithm is sub-optimum in that it suspends more transmissions than necessary to meet the protection requirements. We propose an alternative move-list algorithm that achieves greater efficiency than the standard algorithm at the expense of additional computational complexity. The performance of each algorithm is evaluated with simulated CBSD deployments in the vicinity of federal protection areas.

Section II describes the standard and proposed move-list algorithms. Section III compares the results of each algorithm for three different protection areas. Section IV summarizes the results and draws conclusions.

II. Move-List Algorithms

Given a set of transmissions that overlap in frequency with a protected channel, a move-list algorithm identifies which transmissions must be suspended (and possibly relocated to a different channel) to avoid excessive interference in a protected federal incumbent area. In the parlance of the CBRS specifications, an authorization to transmit is called a "grant." Hence, a move list is a list of grants that must be suspended when a federal incumbent protection area becomes active. Reasons for activation of a protection area on a given channel include detection of a federal incumbent signal within the protection area on that channel.

To obtain the move list, the algorithm computes the path loss from each transmitter to a point in the protected area and, using a stochastic model for the loss on each link, computes the distribution of the aggregate interference at that point. The algorithm then chooses a subset of the grants that must be suspended (relocated) such that the 95th percentile of the aggregate interference is below a threshold at any point in the protected area. Because the protection requirement is based on the 95th percentile of the aggregate interference, it is necessary that all SASs managing CBSDs in the vicinity of a protected area exchange grant information and execute the same move-list algorithm on the total grant population, hence the need for standardization of the move-list algorithm.

This section describes three move-list algorithms considered in the analysis below. The "standard" algorithm is that which is currently in force in CBRS industry standards; the "joint-azimuth" algorithm generates a smaller (i.e., more efficient) move list for a given deployment and protection threshold at the expense of additional computational complexity; and the third algorithm achieves much of the gain of the joint-azimuth algorithm but with the complexity of the standard algorithm.

A. Standard Algorithm

For a given channel, point in the protection area, and protection threshold, the standard move-list algorithm [2] first sorts the grants by their median interference contribution to the protection point. It then identifies the minimal portion of the sorted list, starting with the strongest interferer, that must be removed from the channel so that the 95th percentile of the aggregate interference of the remaining grants does not exceed the protection threshold. The grants that must be removed to meet the protection threshold are placed on the move list.

Because the incumbent receiver antenna is directive in the azimuthal plane and its direction is variable, the algorithm must take into account every possible azimuth direction of the incumbent. For each possible direction, it must apply the azimuthal gains of the transmit and receive antennas accordingly, depending on the bearing of each transmitter relative to the protection point. The standard move-list algorithm accounts for the range of possible azimuths of the incumbent by determining the cutoff for every possible receiver azimuth separately, resulting in a component move list for each azimuth, and taking the union of those component move lists. This process is repeated for every protection point in the protection area, and the move list for the protection area is the union of the move lists of the points.

Pseudocode for the standard move-list algorithm is given in Algorithm 1. Line 5 finds the largest cutoff to the sorted list such that the 95th percentile of the aggregate interference does not exceed the protection threshold. The search for this cutoff can be performed with a binary search.

Algorithm 1: Standard move-list algorithm

Input: Protection channel c, protection threshold t, set of protection points $\mathcal{P}$, set of grants $\mathcal{G}$
Output: Move list for channel c, $\mathcal{M}_c \subseteq \mathcal{G}$

// loop through every protection point
1 **foreach** *point* p *in set* $\mathcal{P}$ **do**
  // find the grants in the "neighborhood" of protection point p *and channel* c
2   $\mathcal{G}_{c,p} \leftarrow$ `Neighborhood`$(\mathcal{G}, c, p)$;   *//* $\mathcal{G}_{c,p} \subseteq \mathcal{G}$
  // sort grants by their median interference contribution, $P_i(c) + G_i(p) - L_i(p)$*, smallest to largest, where* $P_i(c)$ *is the conducted power of the* i*th grant in channel* c *in dB relative to 1 mW (dBm),* $G_i(p)$ *is the transmit antenna gain in the direction of point* p *in dB relative to isotropic (dBi), and* $L_i(p)$ *is the median path loss from the transmitter to point* p *(dB)*
3   $\mathbf{S} \leftarrow$ `Sort`$(\mathcal{G}_{c,p})$;   *//* $\mathbf{S} = [G_1, G_2, \ldots, G_N]$
  // loop through every receiver azimuth
4   **for** $a \leftarrow$ minAzimuth **to** maxAzimuth **do**
5     $n_{\max} \leftarrow$ largest n s.t. `95thPrcntl`$(\{G_1, \ldots, G_n\}, a) \leq t$;
6     $\mathcal{M}_{c,p,a} = \{G_{n_{\max}+1}, \ldots, G_N\}$;
7   $\mathcal{M}_{c,p} = \bigcup_a \mathcal{M}_{c,p,a}$;
8 $\mathcal{M}_c = \bigcup_p \mathcal{M}_{c,p}$;

It is important to note that the median interference calculation used for sorting the grants (Line 3 in Algorithm 1) does not include the receive antenna gain. Hence, it is possible for the component move list of a given azimuth to include a grant that is higher on the sorted list in terms of median interference even though it is outside the main beam of the receive antenna, leading to sub-optimality of the standard algorithm. Advantages of the standard algorithm, on the other hand, are that the sort need only be done once per protection point and the calculation of the per-azimuth component move lists (the for-loop at Line 4) can be parallelized.

B. Joint-Azimuth Algorithm

The joint-azimuth move-list algorithm departs from the standard algorithm in a key respect. Instead of treating each receiver azimuth independently of the others, it finds the receiver azimuth for which the aggregate interference is largest, that is, the worst-case azimuth. Then, starting with the highest contributor in terms of median interference contribution, it adds grants to the move list until the aggregate interference at that azimuth is no longer the highest; that is, until another azimuth is the worst case. It repeats this process until the aggregate interference at every azimuth is no greater than the protection threshold. Pseudocode for the joint-azimuth move-list algorithm is given in Algorithm 2.

Algorithm 2: Joint-azimuth move-list algorithm

Input: Protection channel c, protection threshold t, set of protection points $\mathcal{P}$, set of grants $\mathcal{G}$
Output: Move list for channel c, $\mathcal{M}_c \subseteq \mathcal{G}$

// loop through every protection point
foreach *point* p *in set* $\mathcal{P}$ **do**
  // find the grants in the "neighborhood" of protection point p *and channel* c
  $\mathcal{G}_{c,p} \leftarrow$ `Neighborhood`$(\mathcal{G}, c, p)$;   *//* $\mathcal{G}_{c,p} \subseteq \mathcal{G}$
  $\mathcal{M}_{c,p} \leftarrow \emptyset$;
  while $\max_{a \in \text{Azimuths}}$ `95thPrcntl`$(\mathcal{G}_{c,p} \setminus \mathcal{M}_{c,p}, a) > t$ **do**
    $a_{\max} \leftarrow$ azimuth with highest aggregate interference;
    $a_2 \leftarrow$ azimuth with 2nd highest aggregate interference;
    $I_2 \leftarrow$ `95thPrcntl`$(\mathcal{G}_{c,p} \setminus \mathcal{M}_{c,p}, a_2)$;
    // sort grants by their median interference contribution, $P_i(c) + G_{tx,i}(p) - L_i(p) + G_{rx,i}(p, a_{\max})$*, smallest to largest, where* $P_i(c)$ *is the conducted power of the* i*th grant in channel* c *(dBm),* $G_{tx,i}(p)$ *is the transmit antenna gain in the direction of point* p *(dBi),* $L_i(p)$ *is the median path loss from the transmitter to point* p *(dB), and* $G_{rx,i}(p, a_{\max})$ *is the receive antenna gain in the direction of the transmitter*
    $\mathbf{S} \leftarrow$ `Sort`$(\mathcal{G}_{c,p} \setminus \mathcal{M}_{c,p})$;
    add grants to $\mathcal{M}_{c,p}$ starting with the largest in $\mathbf{S}$ until `95thPrcntl`$(\mathcal{G}_{c,p} \setminus \mathcal{M}_{c,p}, a_{\max}) \leq \max(I_2, t)$;
$\mathcal{M}_c = \bigcup_p \mathcal{M}_{c,p}$;

To compare the complexity of the standard and joint-azimuth algorithms, let $|\mathcal{A}|$ be the number of receiver azimuths, let $|\mathcal{G}|$ be the number of grants, and let $|\mathcal{M}|$ be the size of the move list. The complexity of each algorithm is dominated by the calculation of the 95th percentile of the aggregate interference. The standard algorithm performs this calculation approximately $|\mathcal{A}| \log |\mathcal{G}|$ times, assuming a binary search for Line 5 of Algorithm 1. The joint-azimuth algorithm, on the other hand, performs this calculation approximately $|\mathcal{A}||\mathcal{M}|$ times. Hence, the complexity of the joint-azimuth algorithm grows with $|\mathcal{M}|$ while that of the standard algorithm grows with $\log |\mathcal{G}|$. For large move lists, $|\mathcal{G}|$ and $|\mathcal{M}|$ are of the same order and can be in the thousands or tens of thousands. Furthermore, the standard algorithm can parallelize the processing of both the protection points and the azimuths, while the joint-azimuth algorithm can only parallelize the protection points.

The joint-azimuth algorithm constructs the move list by considering all possible receiver azimuths jointly rather than independently as in the standard algorithm. Further gains could be achieved by jointly considering all protection points as well as all azimuths, but the number of protection points can be in the thousands in practice.

C. *Modified Standard Algorithm*

The modified standard algorithm is a minor variation of the standard algorithm. Instead of sorting grants by their *median* interference contribution at Line 3 of Algorithm 1, the modified standard algorithm sorts them by a higher percentile of their interference contribution. Otherwise, the two algorithms are identical, with no difference in computational complexity. We heuristically found that sorting the grants by the 99th percentile of their interference contribution yields the smallest move lists in the examples of the next section.

III. Analysis

We executed the move-list algorithms described in Section II on simulated deployments of CBSDs along coastal areas of the U.S. The National Telecommunications and Information Administration (NTIA) has defined federal incumbent protection areas along the U.S. coasts [3]. The simulated deployments were conducted in the vicinity of three of these protection areas: along the northwest coast of the continental U.S. near La Push, Washington; along the Gulf coast near Pensacola, Florida; and along the East coast near Daytona Beach, Florida. These areas were chosen to yield small, medium, and large move lists, respectively.

Besides the CBSD deployment, additional inputs to the move-list algorithm, such as the protection threshold and the height and beamwidth of the incumbent receiver antenna, are specific to the area being protected and are given in keyhole markup language (KML) files [4].

The move-list calculations were performed using the reference implementations of the propagation model, antenna models, and the standard move list algorithm [5] and the reference geodata (terrain elevation and land classification) [6] employed in SAS certification testing.

A. *CBSD Deployment Model*

We used a deployment model with the assumptions outlined below to distribute two categories of CBSDs. Category A CBSDs are lower power devices with a maximum effective isotropic radiated power (EIRP) of 30 dBm/10 MHz and are typically installed indoors. Category B CBSDs are higher power devices (47 dBm/10 MHz maximum EIRP) and are professionally installed outdoors [1].

TABLE I
Users Served by Each CBSD Category

Area Type	Percent Users Served by		Users per CBSD	
	Cat. A	Cat. B	Cat. A	Cat. B
Urban	80 %	20 %	50	200
Suburban	60 %	40 %	20	200
Rural	40 %	60 %	3	500

TABLE II
CBSD Antenna Height and EIRP

Area Type	Antenna Height (m)		EIRP (dBm/10 MHz)	
	Cat. A	Cat. B	Cat. A	Cat. B
Dense Urban	50 %: 3 to 15 25 %: 18 to 30 25 %: 33 to 60	6 to 30	26	40 to 47
Urban	50 %: 3 50 %: 6 to 18	6 to 30	26	40 to 47
Suburban	70 %: 3 30 %: 6 to 12	6 to 100	26	47
Rural	80 %: 3 20 %: 6	6 to 100	26	47

The model uses Geographic Information System (GIS) 2011 National Land Cover Database (NLCD) data [7] and 2010 U.S. Census population data [8] to distribute the CBSDs. CBSDs were deployed to each census tract based upon population and land coverage classification. CBSDs were placed as far as 250 km from the protection area boundary for Category A and as far as 600 km from the boundary for Category B.

First, the NLCD area classification codes were grouped and mapped to dense urban, urban, suburban, and rural regions and applied for each census tract. For the dense urban and urban regions, a daytime traveling factor expressed in terms of a percentage was included to account for the higher population densities that occur in cities during daytime [9].

Assuming a mature deployment, a market penetration factor of 20 % was assumed for this band. To account for distribution of users across ten available 10 MHz channels in the 3.5 GHz band, a scaling factor of 10 % was also included to determine the number of effective users potentially operating in each 10 MHz channel of interest.

The percentages of users served by Category A and Category B CBSDs vary depending on the classification regions and are noted in Table I. After calculating the number of users served by each category of CBSD, the numbers of Category A and Category B CBSDs were calculated based upon the effective number of users per CBSD, also shown in Table I.

The antenna height and EIRP of each CBSD was randomized with a uniform distribution based upon the land usage classification as shown in Table II. The antenna was assumed to be omnidirectional.

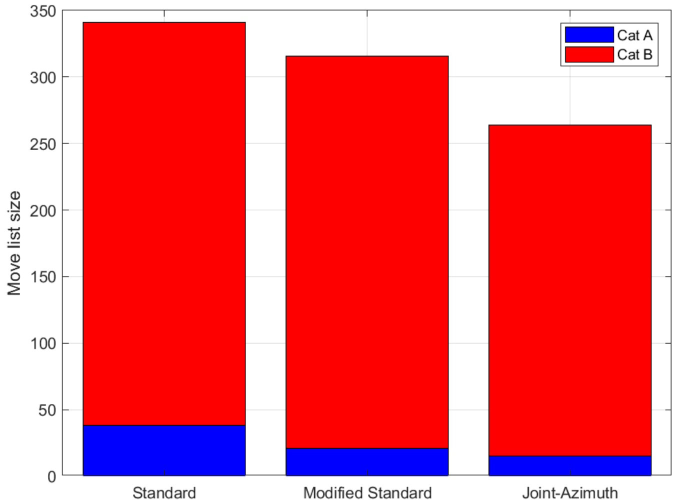

Fig. 1. Move list size, La Push, Washington, offshore protection area.

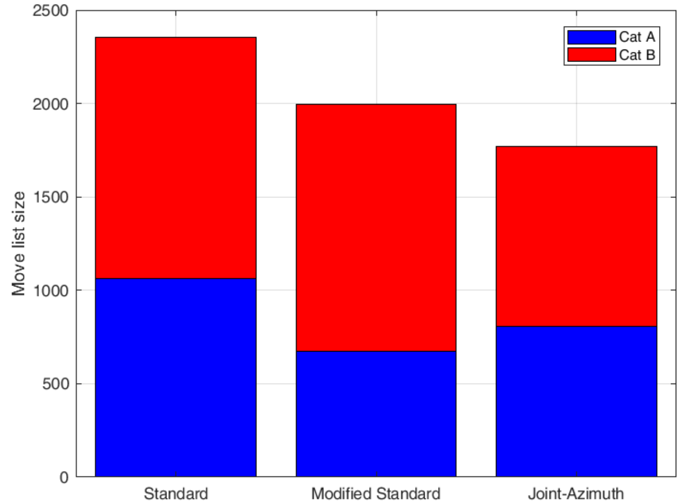

Fig. 2. Move list size, Pensacola, Florida, protection point.

B. Protection Area Examples

1) Small Move List: As an example of a small move list, the algorithms were executed on the La Push coastal protection area. Each move list was calculated for a set of 50 protection points (35 points on the contour and 15 interior points). As shown in Fig. 1, the standard algorithm generates a move list of size 341 grants, 38 of which are from Category A CBSDs and the remainder are from Category B CBSDs. The joint-azimuth algorithm generates a total move list of size 264 grants, for a savings of 23 % over the standard algorithm, while the modified standard move list has 316 grants for a more modest savings of 7 %.

2) Medium Move List: The second example is for a protection area consisting of a single point, a training site in Pensacola, Florida. While the La Push protection area move lists are in the low hundreds, the move lists for Pensacola are in the low thousands. In this "medium" move list example, the efficiency gains of the joint-azimuth and modified standard move lists are somewhat greater than for the small move list. The joint-azimuth move list is 25 % smaller than the standard move list, while the modified standard move list is 15 % smaller than the standard move list (see Fig. 2), achieving more than half of the gain of the joint-azimuth list. However, the efficiency of the modified standard list is obtained solely in Category A grants, while the joint-azimuth list efficiencies are spread across Category A and Category B.

To illustrate how the algorithms utilize the available interference budget, we calculated the 95th percentile of the aggregate interference at the protection point after applying each move list. Fig. 3 plots the 95th percentile as a function of the incumbent receiver azimuth. The horizontal dash-dot line indicates the protection level for Pensacola of −139 dBm/10 MHz. While all three algorithms keep the interference below the protection level, it is clear from this plot that the joint-azimuth algorithm packs transmissions within the interference budget much more efficiently than the other algorithms at almost every azimuth.

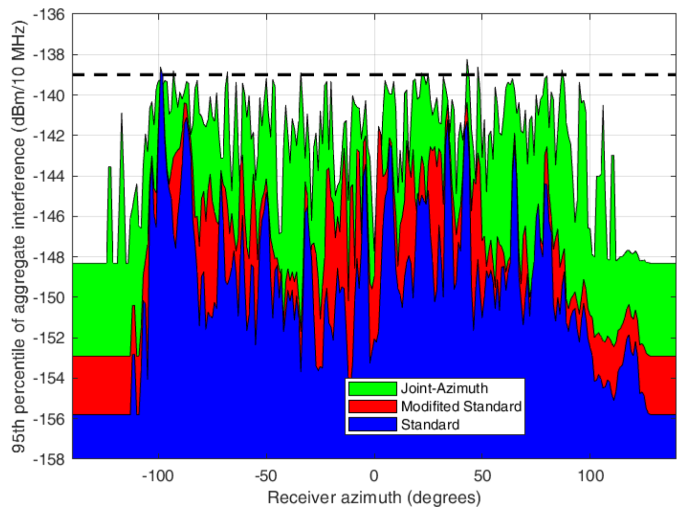

Fig. 3. Aggregate interference, Pensacola, Florida, protection point.

3) Large Move List: The move lists for the Daytona Beach protection area are 3 to 4 times larger than those of Pensacola (see Fig. 4) and were calculated for a set of 18 protection points along the contour of the protection area near the shore. The joint-azimuth move list is 18 % smaller than the standard move list, and the modified standard move list follows closely at 13 % smaller. Here again, the savings of the modified standard list are solely in Category A grants, while the joint-azimuth list has fewer grants in both categories.

Across all three examples, we observe that the performance of the modified standard algorithm relative to that of the joint-azimuth algorithm improves with increasing move list size. Furthermore, the aggregate interference results (Fig. 5) show once again that these algorithms use the interference budget more efficiently than the standard algorithm.

A geographic view of the difference between the standard

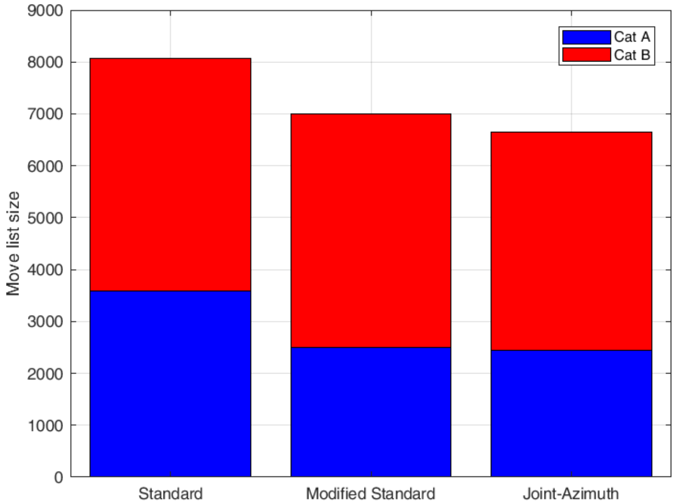

Fig. 4. Move list size, Daytona Beach, Florida, offshore protection area.

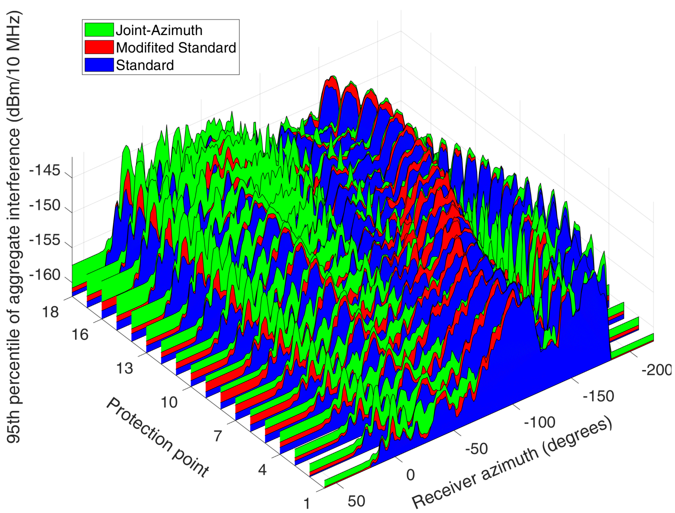

Fig. 5. Maximum aggregate interference, Daytona Beach, Florida, offshore protection area.

and joint-azimuth move lists is shown in Fig. 6.[1] Blue markers indicate CBSDs with grants on the standard move list but not on the joint-azimuth move list. Yellow markers are on the joint-azimuth move list but not on the standard move list (there is only one such CBSD). Markers with dots represent Category B CBSDs, while markers without dots are Category A. We observe that the excess Category A CBSDs on the standard move list are all within 90 km of the protection area, while the excess Category B CBSDs are 160 km to 375 km from the protection area.

[1]Certain commercial products are identified in this paper in order to specify the experimental results adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the products identified are necessarily the best available for the purpose.

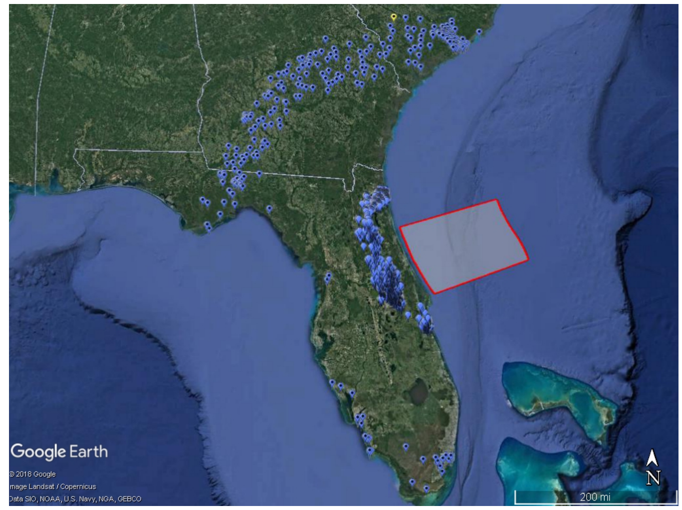

Fig. 6. Standard/joint-azimuth move list difference, Daytona Beach, Florida, offshore protection area. Blue marker: diff on standard move list; yellow marker: diff on joint-azimuth move list; blank marker: Category A; dot marker: Category B.

IV. Conclusion

An analysis of alternative federal incumbent protection algorithms for the 3.5 GHz band finds that a more efficient selection of the CBRS transmissions to be moved from a protected channel results in 18 % to 25 % fewer transmissions being affected compared with the current standardized algorithm. The improved algorithm achieves these gains by considering all possible receiver azimuths jointly rather than independently, but at considerable additional computational cost. Alternatively, a simple change to the standard algorithm achieves much of these gains, mainly for Category A grants, and with no increase in computational complexity.

References

[1] "Citizens broadband radio service," 2 C.F.R. § 96, 2016.
[2] "Requirements for commercial operation in the U.S. 3550–3700 MHz citizens broadband radio service band," Wireless Innovation Forum Document WINNF-TS-0112, Version V1.5.0, May 2018. [Online]. Available: https://workspace.winnforum.org/higherlogic/ws/public/document?document_id=6531
[3] Letter from Paige R. Atkins, Assoc. Admin., Office of Spectrum Mgt., NTIA, to Julius P. Knapp, Chief, Office of Eng. and Tech., FCC, May 2018. [Online]. Available: https://www.fcc.gov/ecfs/filing/1051705880764
[4] Dynamic Protection Area kml files, May 2018. [Online]. Available: https://github.com/Wireless-Innovation-Forum/Spectrum-Access-System/tree/master/data/ntia
[5] SAS Testing and Interoperability Repository. [Online]. Available: https://github.com/Wireless-Innovation-Forum/Spectrum-Access-System
[6] SAS Data Repository. [Online]. Available: https://github.com/Wireless-Innovation-Forum/SAS-Data
[7] 2011 National Land Cover Database, U.S. Geological Survey, 2011. [Online]. Available: http://viewer.nationalmap.gov/basic/
[8] 2010 National Census Tracts Gazetteer, U.S. Census Bureau, 2010. [Online]. Available: http://www.census.gov/geo/maps-data/data/gazetteer2010.html
[9] Commuter Adjusted Daytime Population: 2006-2010, U.S. Census Bureau. [Online]. Available: https://www.census.gov/data/tables/2010/demo/metro-micro/commuting-employment-2010.html